Automated astrometric analysis of satellite observations using wide-field imaging

Jovan Skuljan and John Kay

Defence Technology Agency, Auckland, New Zealand

ABSTRACT

An observational trial was conducted in the South Island of New Zealand from 24 to 28 February 2015, as a collaborative effort between the United Kingdom and New Zealand in the area of space situational awareness. The aim of the trial was to observe a number of satellites in low Earth orbit using wide-field imaging from two separate locations, in order to determine the space trajectory and compare the measurements with the predictions based on the standard two-line elements. This activity was an initial step in building a space situational awareness capability at the Defence Technology Agency of the New Zealand Defence Force. New Zealand has an important strategic position as the last land mass that many satellites selected for deorbiting pass before entering the Earth's atmosphere over the dedicated disposal area in the South Pacific.

A preliminary analysis of the trial data has demonstrated that relatively inexpensive equipment can be used to successfully detect satellites at moderate altitudes. A total of 60 satellite passes were observed over the five nights of observation and about 2600 images were collected. A combination of cooled CCD and standard DSLR cameras were used, with a selection of lenses between 17 mm and 50 mm in focal length, covering a relatively wide field of view of 25 to 60 degrees. The CCD cameras were equipped with custom-made GPS modules to record the time of exposure with a high accuracy of one millisecond, or better. Specialised software has been developed for automated astrometric analysis of the trial data. The astrometric solution is obtained as a two-dimensional least-squares polynomial fit to the measured pixel positions of a large number of stars (typically 1000) detected across the image. The star identification is fully automated and works well for all camera-lens combinations used in the trial. A moderate polynomial degree of 3 to 5 is selected to take into account any image distortions introduced by the lens. A typical RMS error of the least-squares fit is about 0.1 pixels, which corresponds to about 4 to 10 seconds of arc in the sky, depending on the pixel scale (field of view). This gives a typical uncertainty between 10 and 25 metres in measuring the position of a satellite at a characteristic range of 500 kilometres.

The results of this trial have confirmed that wide-field measurements based on standard photographic equipment and using automated astrometric analysis techniques can be used to improve the current orbital models of satellites in low Earth orbit.

1. INTRODUCTION

A new space situational awareness (SSA) capability has been built in the Defence Technology Agency (DTA) over the past couple of years as a result of collaboration between New Zealand and the United Kingdom. In early 2014, DTA was contacted by the UK's Defence Science and Technology Laboratory (Dstl) with a proposal to jointly observe the shallow atmospheric re-entry of the Automated Transfer Vehicle 5 (ATV-5), launched by the European Space Agency (ESA) and scheduled for deorbiting in February 2015. According to the original re-entry plans by ESA, the South Island of New Zealand was the last large land mass that the satellite would pass, at an average altitude of 100 km, before plunging into the South Pacific Ocean. The proposed trajectory profile of ATV-5 differed significantly from a typical steep descent that was common for other satellite re-entries in the past. A shallow profile was chosen in this case, to gain some additional knowledge about the behaviour of a satellite in such a scenario. This provided an opportunity to observe the event from the ground and collect valuable scientific data. Unfortunately, the original plans for a shallow re-entry of ATV-5 had to be changed by ESA due to an on-board failure. The ATV-5 was sent to burn in the Earth's atmosphere on a steep descent instead, and the event was not visible from New Zealand.

In spite of the change of plan with the ATV-5 re-entry, a decision was made to proceed with the joint NZ-UK observing trial in the South Island of New Zealand. A new list of targets was prepared, mainly including satellites (or rocket bodies) in low Earth orbit (LEO), with uncertain orbital parameters. The list was expanded with some additional satellites with well-known orbits, to provide a test case for the astrometric calibration.

2. OBSERVATIONS

The observation trial was conducted between 24 and 28 February 2015 from two separate locations in the South Island, one at the Mt John Observatory, Lake Tekapo, and the other at the National Institute of Water and Atmospheric Research (NIWA) station at Lauder, Central Otago. The straight-line distance between the two locations is about 130 km, which offers a solid base for triangulation. When the same satellite is observed from both sites, it is possible to determine the location of the satellite in three dimensions. However, the triangulation analysis of images recorded during this trial is still a work in progress and will not be discussed any further in this paper.

A selection of scientific CCD detectors and ordinary DSLR cameras were used during this trial, combined with a set of standard medium to wide field lenses. A typical field of view was between 25° and 60°, depending on the detector size and the focal length of the optics. Due to the relatively large field of view, most images showed increased geometrical distortions towards the edges. These distortions were treated as part of the coordinate transformation during the astrometric calibration by simply increasing the polynomial degree of the least-squares fit.

The complete list of equipment used during the trial is shown in Table 1. Some of the cameras were mounted on motorised equatorial mounts to compensate for the Earth's rotation. This enabled an easy subtraction of the stellar background using a number of frames taken before and after the image containing a satellite trace, for an increased detection probability, especially in the case of faint satellites.

Table 1. List of equipment

Location	Camera	Lens	Mount
Mt John	Canon EOS 1000D	Canon EF 50mm f/1.4	Sky-Watcher EQ6 Pro
	Canon EOS 1200D	Canon EF 50mm f/1.8 II	Fixed tripod
	QSI 640ws	Canon EF 24mm f/1.4L II USM	Sky-Watcher EQ6 Pro
NIWA Lauder	Canon EOS 40D	Canon EF 50mm f/1.8 II	Sky-Watcher EQ6 Pro
	Canon EOS 6D	Canon EF 50mm f/1.8 II	Sky-Watcher EQ6 Pro
	QSI 640ws	Canon EF 24mm f/1.4L II USM	Sky-Watcher EQ6 Pro
	Olympus E-M1	Olympus M.Zuiko 17mm f/1.8	Fixed tripod

A total of 2606 images were collected during this trial. About 8% of these images were classified as test frames, while another 8% were dark images. The remaining 84% (2176 images) were target frames, each showing a satellite trace recorded against the stellar background. A total of 60 satellite passes were observed. In some cases, the same satellite was recorded on multiple nights. The total number of different satellites is 44, most of them belonging to the LEO group, with a median altitude of about 720 km at perigee and 780 km at apogee.

3. DATA REDUCTION

A typical image of a satellite trace against the stellar background is shown in Fig. 1. The image was recorded on 26 February 2015 at 10:32:16 UTC from the Mt John Observatory site using the QSI 640ws CCD camera and Canon EF 24mm lens. The satellite shown in the image is Fengyun 3A (NORAD 32958) passing through the constellation of Dorado, just north of the Large Magellanic Cloud. Some of the brightest stars are labelled for reference. The CCD detector resolution is 2048 x 2048 pixels and each pixel is a 7.4 µm square. The nominal image scale is about 62 arc seconds per pixel, while the field of view is about 34.5° on each side (48.8° diagonal). The total exposure duration

was about 6 seconds, but the mechanical shutter was repeatedly closed and opened during the exposure to insert some gaps in the satellite trace, for an increased number of reference points. This technique has been suggested as a possible solution in a situation when the readout time is relatively long and taking separate images at higher frame rate is not possible [1].
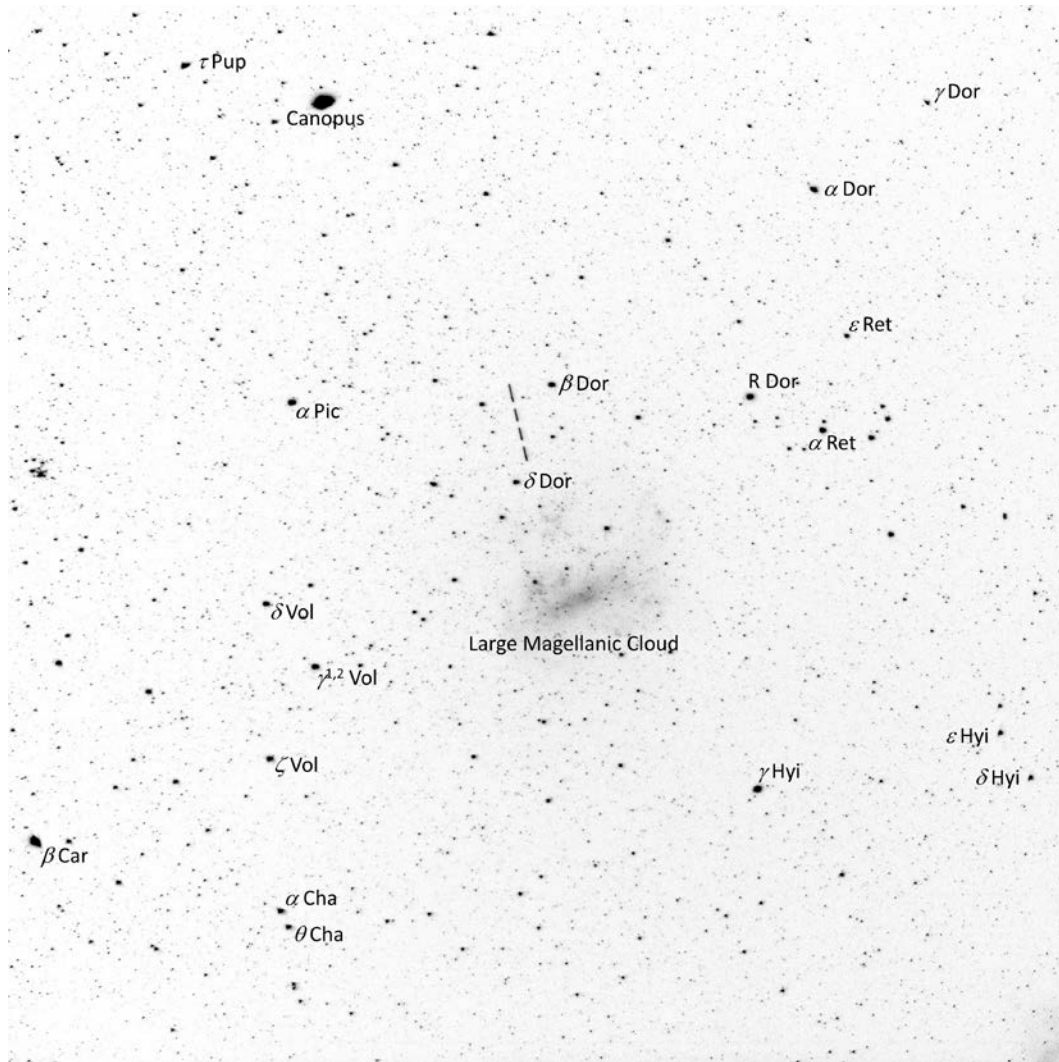

Fig. 1. Fengyun 3A recorded on 26 February 2015 from Mt John. North is up. The declination range covered by the image is from -50° to -84°.

4. ASTROMETRIC CALIBRATION

A variety of different algorithms for astrometric calibration of large sets of astronomical images are described in the literature [2]. Most of these are specific to a certain instrument, or a certain project, and cannot be easily applied to a different optical setup. There exists, however, a much more robust algorithm, known as *Astrometry.net* [3], which we originally wanted to include in our data analysis. Unfortunately, we found that the Astrometry.net procedures did not perform as expected on our images and a decision was made to search for a different algorithm. As a result, a specialised software tool called *StarView* was developed. The software was written in Delphi and has a rich graphical user interface with a variety of features including a display of the FITS (Flexible Image Transport System)

image, adjustment of the contrast, colour map and zoom level, examination of individual pixel values, image histogram and basic statistics. An example of the image display is shown in Fig. 2.
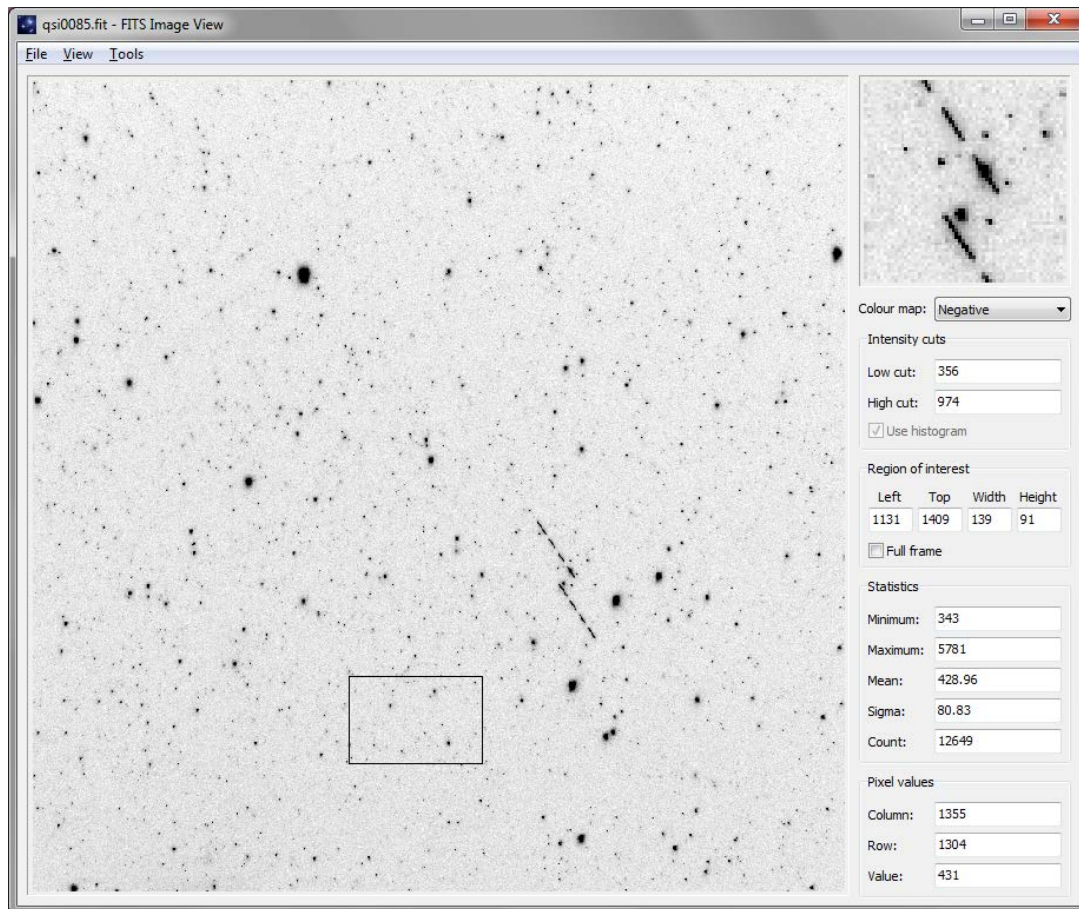

Fig 2. FITS image display in StarView

StarView has its own data base of 117955 stars taken from the new reduction of the Hipparcos catalogue [4]. The Hipparcos proper motions are combined with stellar radial velocities for over 35000 stars from the Pulkovo Radial Velocity Catalogue, for an improved space motion model during the transformation between the catalogue epoch and time of observation. In addition, the Bright Star Catalogue is used for the identification of bright stars in the calibrated images. All catalogue data were obtained from the SIMBAD Astronomical Database in Strasbourg, France [5].

A star chart can be displayed in StarView at a suitable orientation and pixel scale, to match any given FITS image, or a larger area around the image, for a preliminary identification of the sky region. A typical star chart display is shown in Fig 3.

The most important feature of StarView is that it is capable of automatically performing a completely blind calibration of any given image, within a certain range of parameters. The software has been tested on images taken with a variety of cameras and lenses, as listed in Table 1 above. The only requirement is that the nominal pixel scale (in arc seconds per pixel) is specified as an input parameter. This is simply obtained from the physical size of one pixel on the detector and the focal length of the lens. The star identification algorithm will work even if the nominal pixel scale is known to a limited accuracy of a few per cent. Once the astrometric calibration is complete, the actual (true) pixel scale is automatically calculated and included in the final list of output parameters.
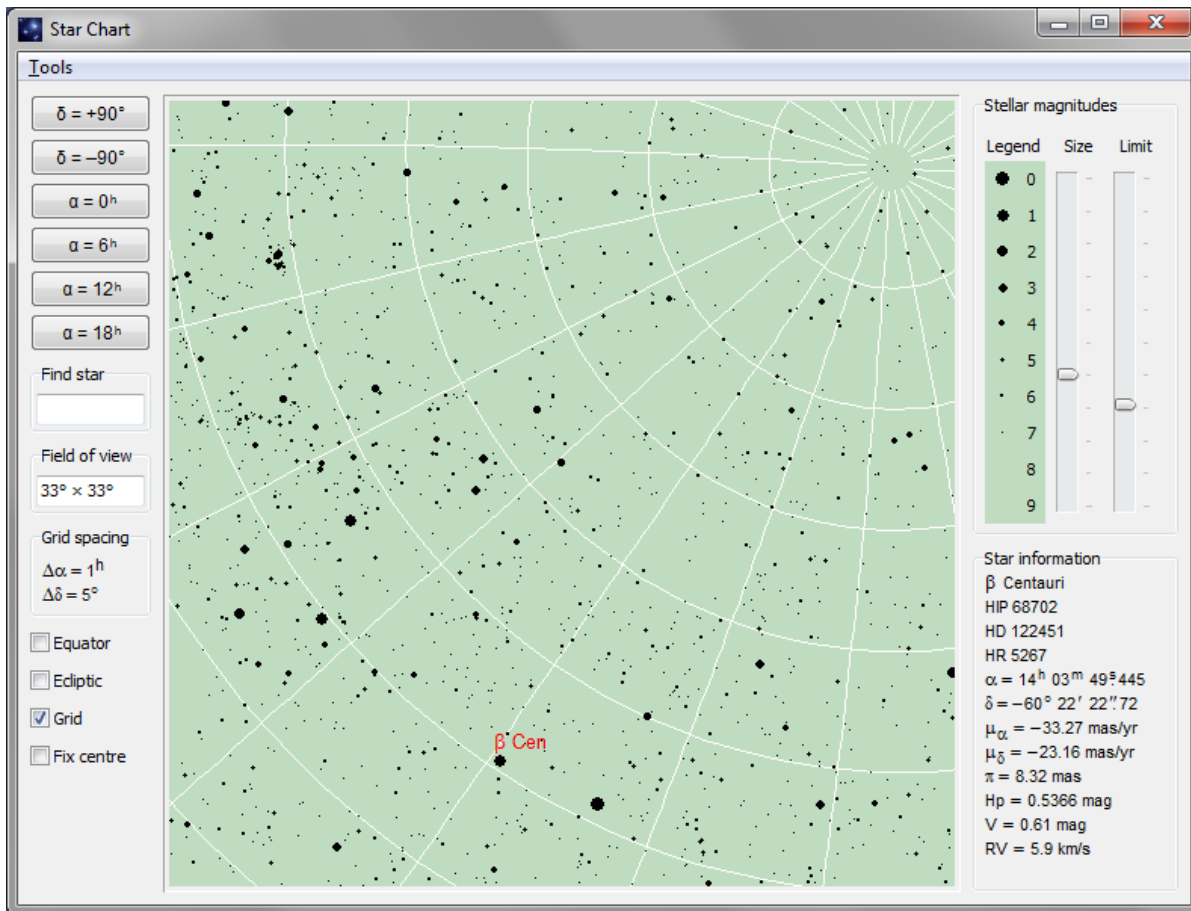
Fig. 3. Star chart display in StarView

A typical astrometric calibration for a given image, as performed fully automatically by StarView, consists of the following steps:

1. A two-dimensional wavelet transform is applied to the raw frame using a so-called "Mexican hat" profile. This is done in order to eliminate any unwanted scattered background light and enhance the stellar objects and satellite traces. The fast Fourier transform technique is applied, using the FFTW (The Fastest Fourier Transform in the West) library [6].

2. The filtered image is examined pixel by pixel to detect the stars above a given threshold level. At this stage, a star is defined as a reasonably round group of connected pixels not exceeding a specified maximum diameter. For any such group, the approximate coordinates of the centre are calculated by averaging the X and Y pixel positions. The threshold level is expressed in terms of the standard deviation in the background. A typical value of 50 is used in our images to ensure a successful detection of about 1000-1500 stars.

3. With the approximate positions of stars detected in the previous step, a Gaussian profile is fitted to every star in both X and Y directions. This is done on the unfiltered (raw) image, in order to preserve the original point-spread function. At this stage, any stars with poor Gaussian fits are rejected.

4. A special algorithm is then applied to the fitted pixel positions of stars detected in the image in order to identify the sky region and generate a list of *reference stars*. We have developed a new star identification method called "stellar fingerprints" which will be explained in the next section. The total number of reference stars will depend on the sky region and is usually about 20-40. The minimum number of

reference stars that are required for the remaining parts of the algorithm is 3 and they can be selected manually, if desired. However, the automated process proved to be practically 100% reliable, assuming that the input parameters for a given camera-lens combination are prepared correctly.

5. The initial set of reference stars provides a first approximation for the astrometric solution in the form of a two-dimensional polynomial regression (least-squares fit). For each star, the equatorial coordinates (α, δ) are first projected onto the viewing plane to obtain the projected coordinates (ξ, η) which are still in angular units. A set of polynomial regressions is then applied to convert between the pixel coordinates (x, y) and projected coordinates (ξ, η). The initial polynomial degrees are kept low (1-2) until the total number of stars becomes large enough to allow for higher degrees.

6. Out of all stars detected in the image, a smaller subset of so-called *core* stars is created first, so that the stars are roughly evenly distributed across the frame. In order to achieve this, the image is divided into a number of smaller regions and only a small number of stars are picked from each region. Typically, there will be 200-300 core stars evenly distributed across the image. These core stars will be identified first, based on the initial astrometric solution obtained from the reference stars, as described in the previous step. An iterative process is applied, where only one new core star is added to the solution at a time. The existing solution is first used to predict the equatorial coordinates of the closest unidentified core star and this is then cross-referenced with the star catalogue. If the identification is successful, the star is added to the solution and the process is repeated.

7. The improved astrometric solution based on the core stars is then applied in a single step to predict the equatorial coordinates of the remaining few hundred stars in the image. The final solution will be based on all successfully identified stars. A typical polynomial degree is 4-5 and the RMS error is around 0.1-0.2 pixels (6-12 arc seconds at 60 seconds per pixel).

5. STELLAR FINGERPRINTS

StarView uses its own algorithm for star identification called "stellar fingerprints". A stellar fingerprint is simply a small circle in the sky centred at a given catalogue star and containing all other stars in the immediate neighbourhood. An example is shown in Fig. 4 for the bright star α Pictoris (visual magnitude 3.2). The radius of the circle is chosen carefully, depending on the total field of view covered by a typical CCD image. A stellar fingerprint should not be larger than a fraction of the entire image. In our case, a moderately large radius of 5° has been adopted, so that the circle covers about a quarter of a typical image.

The total number of fingerprints, as well as the number of field stars in each fingerprint, will depend on the limiting stellar magnitude, which is selected based on the average exposure time. Although the majority of our images show stars down to a magnitude of about 9-10, it is impractical to include all these stars in the astrometric analysis, partly because the computing time increases rapidly when more stars are added. We have found by experimenting that at a limiting magnitude of about 7.5-8.0 there is a reasonable compromise between the number of stars used for the calibration and the computing time spent on star identification. Our current data base contains 23270 fingerprints for stars brighter than 7.5. The magnitude limit for the field stars within each fingerprint is somewhat higher than the limit for the central star and is set to 8.0. The total number of field stars in one fingerprint depends on the sky region and ranges between 24 and 213.

For each stellar fingerprint, two additional one-dimensional images are calculated, one showing the position angles of all stars around the central star, which we call the "theta-comb", and the other showing the distances from the central star (the "rho-comb"). The distances are on a logarithmic scale, so that a pair of rho-combs that represent the same stellar region at two slightly different scales can be cross-correlated in order to determine the scaling factor.

With a list of pixel coordinates of the stars detected in a given image, the automatic astrometric calibration algorithm will start by selecting a relatively bright star not far from the image centre. It is very likely that this will be one of

the catalogue stars, assuming that the initial parameters are set correctly. If it happens that the star is not in the catalogue, so that the rest of the algorithm fails, a new star will be selected until a successful identification is achieved. An "observed" stellar fingerprint is created for the selected star, based on the measured pixel coordinates and using the initial estimate for the nominal pixel scale. The observed theta-comb is then cross-correlated with the catalogue theta-combs for every catalogue star. This will generate a list of possible rotation angles for the selected stellar region in the image. At the same time, the observed rho-comb is cross-correlated with the catalogue rho-combs to determine the scaling factors. The rotation angles and scaling factors are used to transform the observed pixel positions and match them with the catalogue positions. Out of all fingerprints examined, the one with the most stars that were successfully identified is selected. The field stars inside the fingerprint become the reference stars for the remaining steps of the astrometric calibration algorithm described in the previous section. The entire process of fingerprint identification takes only a few seconds of computing time.
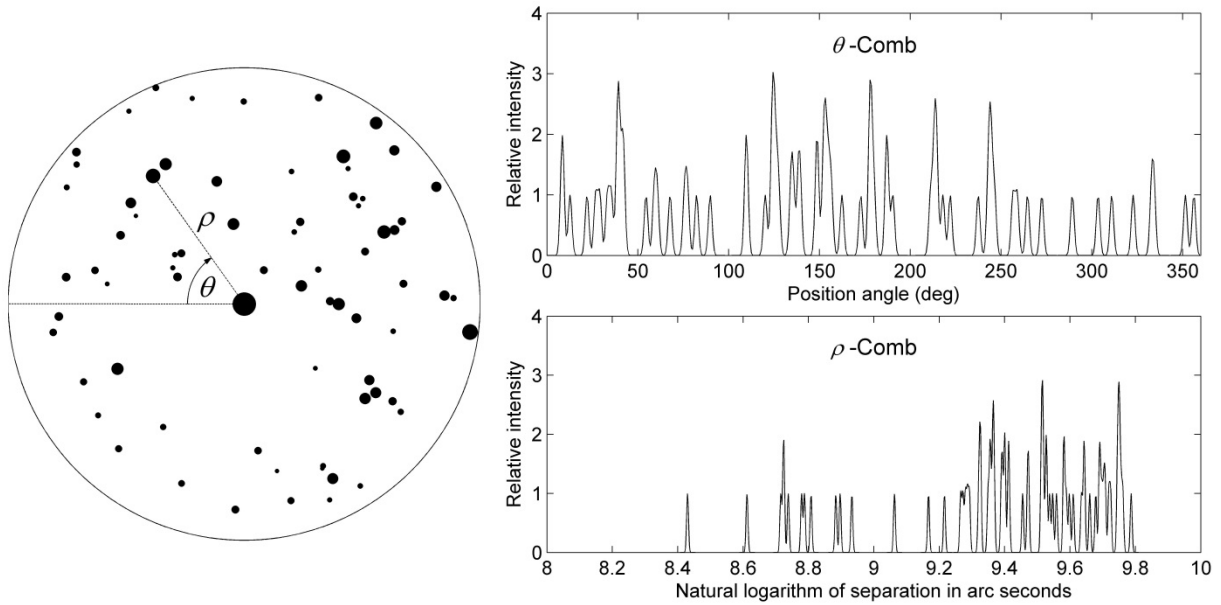

Fig 4. Stellar fingerprint for the bright star α Pictoris

6. SATELLITE MEASUREMENTS

In order to illustrate the astrometric calibration method described in this paper, a number of satellite images were analysed and the equatorial positions of the satellites captured in these images were computed. At this stage, the extraction and analysis of satellite traces is still not part of StarView. We perform these tasks separately using Matlab and ESO-MIDAS (a software tool designed for astronomical images). A total of 37 images of 9 different satellites were analysed. The list of satellites is shown in Table 2, including the satellite name, NORAD catalogue number, COSPAR identification, orbit size (minimum and maximum altitude above the ellipsoid), orbital inclination with respect to the equator, epoch of the orbital elements and the mean age of the elements, in days, with respect to the time of observation. All observations were taken from the Mt John Observatory site on 26 February 2015 using the QSI 640ws CCD camera and the canon 24 mm lens.

The satellite traces were extracted from the images by first fitting a set of one-dimensional Gaussian profiles, either horizontally, or vertically (depending on the tilt of the satellite trace) in order to get the central pixel position along the trace. A straight line was then fitted to the central points to obtain the mean path of the satellite in pixel units. This line was then used to obtain the intensity profile along the satellite trace by summing up a number of pixels on either side (horizontally or vertically). The intensity profile is typically characterised by a number of rising and falling edges. When a new exposure is started, a rising edge is created from the background level to the maximum

intensity. When the exposure is finished, the profile shows a falling edge back to the background level. There might be several rising and falling edges in between, if the mechanical shutter is manipulated during the exposure. However, these additional edges were not analysed for this paper. For each satellite, the two end points on the trace corresponding to the start and end of exposure were accurately measured by fitting a suitable edge profile (a standard Gauss error function, i.e. the integral of the Gaussian profile) and the results were compared with the orbital predictions. The total number of data points extracted from all satellite traces is 104.

Table 2. List of satellites observed from the Mt John Observatory site on 26 February 2015.

Name	NORAD	COSPAR	Orbit size	Incl.	TLE epoch (UTC)	Age(d)
Meteor 1-4 Rocket	4394	1970-037-B	460×522 km	81.2°	26-Feb-2015 21:12	0.48
H-2A R/B	36105	2009-066-B	512×556 km	97.4°	28-Feb-2015 19:24	2.40
Envisat	27386	2002-009-A	765×766 km	98.3°	26-Feb-2015 03:49	-0.27
Fengyun 3A	32958	2008-026-A	820×834 km	98.5°	26-Feb-2015 04:00	-0.27
USA 238 DEB	38773	2012-048-P	1077×1136 km	63.4°	26-May-2015 21:48	89.46
USA 238	38758	2012-048-A	1077×1136 km	63.4°	26-May-2015 21:48	89.46
Cosmos 1833	17589	1987-027-A	837×862 km	70.9°	26-Feb-2015 17:17	0.26
NOSS 3-3 (A)	28537	2005-004-A	956×1257 km	63.4°	02-Mar-2015 05:21	3.74
NOSS 3-3 (C)	28541	2005-004-C	953×1260 km	63.4°	02-Mar-2015 05:21	3.74

The orbital predictions for each satellite were computed using TheSkyX Professional Edition astronomy software by Software Bisque. The input TLE data file was prepared using the best orbital element sets available for each satellite on the Space-Track archive [7]. For the majority of the satellites, the epoch of the elements is within a few days of the time of observation, which ensures that the orbital prediction is correct. However, in the case of the satellite pair USB 238 and USB 238 DEB, the best orbital elements available were over 89 days away from the date of observation, which lead to an extremely unreliable prediction.

The difference between the measured and predicted equatorial position, called the prediction error in this paper, for each satellite observation was obtained as two separate components, shown in Fig. 5.
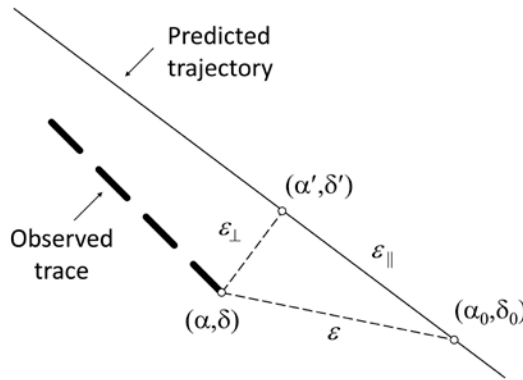

Fig 5. Orbital prediction error components

For a given measured position (α,δ) on the observed satellite trace, we first determine the closest point (α',δ') on the predicted trajectory. We define the *lateral error* $(\varepsilon_\perp)$ as the distance between the measured position (α,δ) and the closest point (α',δ'). The *longitudinal error* $(\varepsilon_\parallel)$ is the distance between the closest point (α',δ') and the predicted position (α_0,δ_0) for the time of observation. The total prediction error is then obtained from $\varepsilon^2 = \varepsilon_\parallel^2 + \varepsilon_\perp^2$. The lateral error is usually determined with more certainty, as the distance between the observed and predicted trajectories can be measured down to the precision of the astrometric solution. On the other hand, the uncertainty in the longitudinal

error can be higher, as it also contains two additional sources of error: inaccurate timing and incorrect placement of the point on the observed trace where the profile intensity reaches exactly 50% of the maximum. In our case, the timing is not an issue, as both the start and end of an exposure are recorded to a high accuracy of better than 1 ms. For a satellite in low Earth orbit, this leads to a positional uncertainty of only a few metres, which is less than our typical astrometric calibration error. However, the determination of the profile edge on the observed satellite trace can be uncertain if the signal to noise ratio is not high enough. No additional analysis was made to determine the longitudinal position uncertainty as a function of the satellite brightness, but it was noticed that for very faint satellites, the error can be up about half a pixel, which is 4-5 times higher than the astrometric RMS error.

The distribution of prediction errors for the nine satellites included in this analysis is shown in Figs. 6 & 7, in angular units (arc seconds) and in kilometres at the satellite position, respectively.
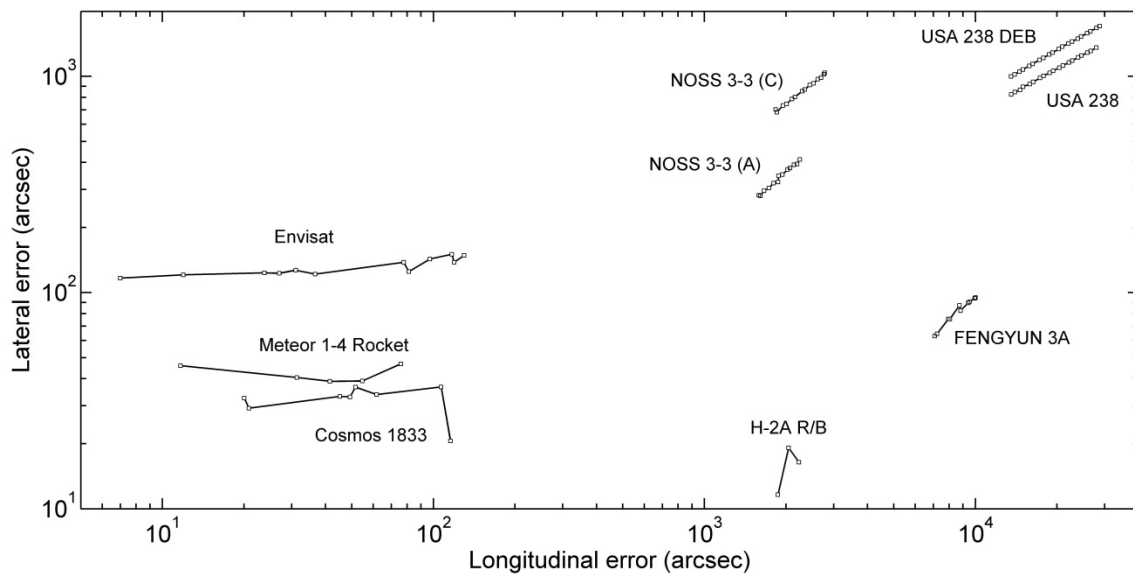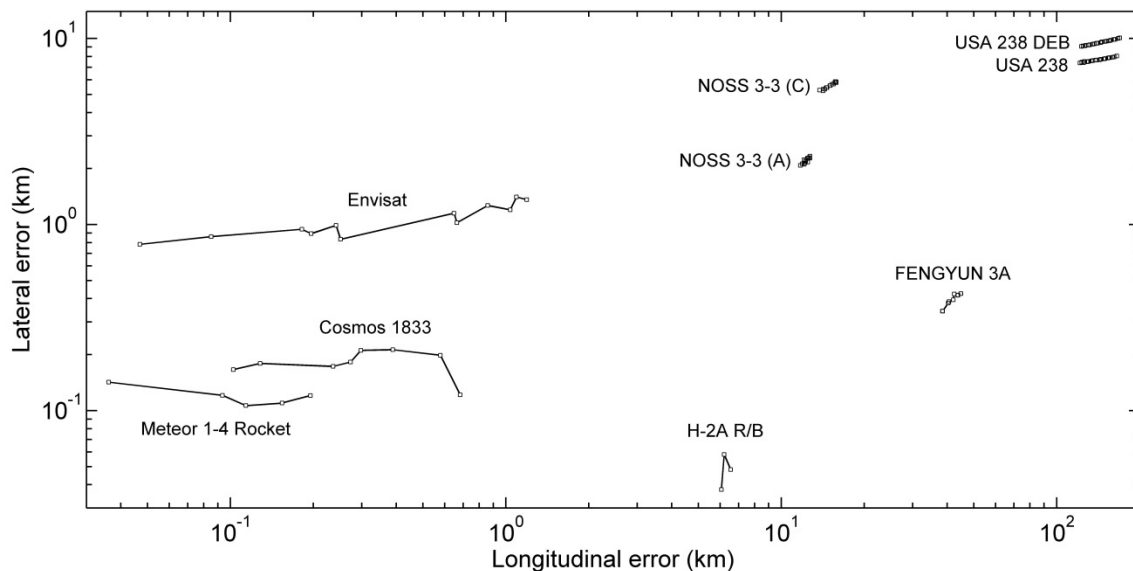

Fig 6. Angular prediction error

Fig 7. Positional prediction error

Both the longitudinal and lateral errors are plotted on a logarithmic scale in order to show all satellites reasonably well, due to the large range of values spanning several orders of magnitude. The angular prediction error can be easily converted to the positional error (in kilometres) when the distance (range) to the satellite is known. The actual range cannot be determined from the observations presented here, as they are made only from one location. However, the *predicted* satellite range obtained from the TLE elements can be used instead for a reliable estimate of the positional error (positional error = angular error × predicted range). We can see that the positional error ranges widely from about 100-200 m in the case of Meteor 1-4 Rocket to 10-20 km for the NOSS 3-3 pair and even 100-200 km (longitudinal error) for the USA 238 pair. In some cases, it is easy to explain such large errors by the fact that the mean epoch of the orbital elements used for the prediction is too far from the observation date, so that the satellite orbit is not well defined. This is certainly the case with the USA 238 pair, as their orbital elements are for an epoch which is exactly three months away from the observation date. Similarly, the NOSS 3-3 pair and H-2A R/B have orbital elements outdated by 2-4 days, which may contribute to their relatively large prediction errors. An interesting outlier is Fengyun 3A, with a longitudinal error of about 40 km, at an approximate range of 1000 km. The TLE epoch for this satellite is only about seven hours away from the observation date. This means that some of the orbital parameters for Fengyun 3A were incorrect.

7.   SUMMARY

A number of satellites from the LEO group were observed during a joint NZ-UK trial with the aim of measuring the position in space and comparing the results with the TLE-based predictions. Some of these observations were analysed and presented in this paper.

A specialised software package called StarView was developed at the Defence Technology Agency for the analysis of wide-filed images covering between 20°-50°. The astrometric calibration is fully automated and is capable of blind identification of the sky region captured in the image. The algorithm uses a new concept of 'stellar fingerprints" to compare the distribution of stars detected in the image with the data base of stars from Hipparcos catalogue. In average, about 1000-1500 stars are successfully identified and used to compute the astrometric solution for an entire image, with a typical RMS error of about 0.1-0.2 pixels.

Some of the satellites observed during the trial were analysed to illustrate this technique. It was demonstrated that relatively inexpensive equipment can be used to successfully locate a satellite in its orbit and compare the results with the orbital predictions.

8.   REFERENCES

1. Privett, G., Timed satellite observations for CCD imagers, *J. Br. Astron. Soc.* 124, 2, 92-95, 2014.
2. Masias, M., Freixenet, J., LLadó, X., Peracaula, M., A review of source detection approaches in astronomical images, *Mon. Not. R. Astron. Soc.* 422, 1674-1689, 2012.
3. Lang, D., Hogg, D. W., Mierle, K., Blanton, M., & Roweis, S., Astrometry.net: Blind astrometric calibration of arbitrary astronomical images, *The Astronomical Journal* 139, 1782–1800, 2010.
4. van Leeuwen, F., Fantino, E, A new reduction of the raw Hipparcos data, Astron. Astrophys. 439, 791-803, 2005.
5. SIMBAD Astronomical Database (online data). Retrieved from: http://simbad.u-strasbg.fr/simbad/
6. Frigo, M., Johnson, S. G., The fastest Fourier transform in the West, Manual for version 3.3.3., 2012. Retreived from: http://www.fftw.org
7. Space-Track, Historical TLE search (online data). Retrieved from: https://www.space-track.org, 23 June 2015.